

Locally stationary wavelet fields with application to the modelling and analysis of image texture

Idris A. Eckley¹

Lancaster University, U.K.

Guy P. Nason

University of Bristol, U.K.

Robert L. Treloar

Unilever Research Ltd.

Summary.

This article proposes the modelling and analysis of image texture using an extension of a locally stationary wavelet process model into two-dimensions for lattice processes. Such a model permits construction of estimates of a spatially localized spectrum and localized autocovariance which can be used to characterize texture in a multiscale and spatially adaptive way. We provide the necessary theoretical support to show that our two-dimensional extension is properly defined and has the proper statistical convergence properties.

Our use of a statistical model permits us to identify, and correct for, a bias in established texture measures based on non-decimated wavelet techniques. The proposed method performs nearly as well as optimal Fourier techniques on stationary textures and outperforms them in non-stationary situations. We illustrate our techniques using piled fabric data from a fabric care experiment and simulated tile data.

Keywords: random field; local spectrum; local autocovariance; texture classification; texture model; non-decimated wavelets

1 Introduction

Wavelet techniques have recently become extremely popular in the statistical literature for nonparametric curve estimation and for the modelling and analysis of time series. For a general overview of wavelet techniques in statistics see the review by Abramovich, Bailey and Sapatinas (2000), Vidakovic (1999) or Nason (2008). This article tackles the problem of modelling and analysing image texture (or more generally, the spatial covariance structure of lattice processes).

Our texture model is based on the locally stationary wavelet (LSW) process model for time series from Nason, von Sachs and Kroisandt (2000) (henceforth NvSK) and we draw our notation largely from this work. Sections 2 and 3 extend the LSW model

¹*Address for correspondence:* Department of Mathematics and Statistics, Lancaster University, Lancaster, LA1 4YF, United Kingdom
E-mail: I.Eckley@lancaster.ac.uk

into two dimensions providing extended versions of NvSK’s model, evolutionary wavelet spectrum, localized autocovariance and their estimators. In many cases the extension is straightforward although in a few cases some non-trivial work is required to check that the necessary theory is still valid. Our new two-dimensional (2D) model provides localized spectra and autocovariance for 2D lattice processes. The localization is incredibly important for applications because the statistical properties of real-life objects often vary with location. For example, Section 5 demonstrates how our model can deal with non-stationary texture classification on simulated tile data.

There are many potential stochastic models for texture. However our key thesis is that texture often has a locally stationary character. Two recent theoretical developments for locally stationary processes are the locally stationary Fourier (LSF) framework (due to Dahlhaus (1997)) and the LSW (due to NvSK). Here we choose LSW because one frequently highlighted aspect of texture is that it possesses structure on many different scales. Moreover, several researchers have highlighted that the human and mammalian visual systems process images in a multiscale manner, preserving both local and global information (see for example, Daugman (1990) or Field (1999)). Thus there is a compelling argument for the development of a multiscale texture model. We are by no means the first to notice this multiscale phenomena or indeed use wavelet techniques in this area and hence provide a synopsis of this field in Section 4.2.

There appears to be no canonical mathematical definition of “texture” although there are plenty of qualitative descriptions. Broadly speaking, texture is the visual character of an image region whose structure is, in some sense, regular: for example the appearance of a woven material. The advent of computational and imaging technology has seen a truly enormous body of work appear on texture. Much of this work focusses on discrimination, classification and segmentation tasks. Section 4 attempts to provide an introduction to the texture modelling and analysis literature.

One of the advantages of possessing a statistical *model* is that its properties can be rigorously defined and discerned. With our model it can be seen that raw use of the popular non-decimated wavelet transform for texture classification (or its variance) is not suitable because, viewed as a spectral quantity, power is inappropriately spread amongst scales and directions. Our statistical theory (and that of NvSK in 1D) shows that this can be ameliorated by a bias correction. Once applied, our method with its bias correction gives superior classification performance compared to the established non-decimated wavelet methods that lack underlying models, see Section 5 for further details.

2 Locally stationary wavelet fields

2.1 Motivation

Suppose we have a random field defined on a regular grid, $\{X_{\mathbf{r}}\}_{\mathbf{r} \in \mathbb{Z}^2}$ for which we wish to estimate the covariance $\text{Cov}(X_{\mathbf{r}}, X_{\mathbf{s}}) = \gamma_{\mathbf{r}, \mathbf{s}}$, where $\mathbf{r}, \mathbf{s} \in \mathbb{Z}^2$. The covariance structure of such a field could take many possible forms. For example, the process could be (second-order) stationary, or intrinsically stationary (see Priestley (1981) or Cressie (1991)) or in

extreme cases the covariance could possess minimal controls such as $\gamma_{\mathbf{r},\mathbf{s}} = \gamma_{\mathbf{t},\mathbf{u}}$ if and only if $\mathbf{r} = \mathbf{t}$ and $\mathbf{s} = \mathbf{u}$, which would permit a high degree of nonstationarity, causing problems for estimation as information about $\gamma(\mathbf{r}, \mathbf{s})$ only comes from x_r and x_s .

The form of covariance structure that we assume lies between the two extremes of stationarity/highly nonstationary form. We permit the covariance structure to change slowly as a function of location. Hence the covariance structure around a particular location, \mathbf{r} , may be estimated by pooling information from data close by. Fields which exhibit this slowly varying structure are termed *locally stationary random fields*. Many real-life images have a locally stationary structure operating at several scales, hence our adoption of wavelets later.

There have been a several developments in the modelling of non second-order stationary spatial processes. See, e.g., Haas (1990), Sampson and Guttorp (1992), Loader and Switzer (1992), Le and Zidek (1992), Le, Sun and Zidek (1997), Higdon, Swall and Kern (1999), Damian, Sampson and Guttorp (2003), Schmidt and O'Hagan (2003). These approaches have predominantly been designed to work with multiple realizations and have considered the more general problem of spatial processes not defined on a regular grid. In particular, these approaches are not multiscale, an important feature in the analysis of textured images (see Section 4.2 for further details).

More recently, wavelet models of the second-order structure were proposed by Nychka, Wikle and Royle (2002) and Mondal and Percival (2008). Nychka *et al.*'s approach for estimating the spatial field covariance structure uses temporal replication to estimate sample covariances (i.e. multiple realizations). Mondal and Percival require a single realization and focus on wavelet variance applied to *stationary* random fields.

Our model, defined below, is distinct from earlier work as it is lattice based, multiscale, permits a locally stationary covariance structure and, critically for texture analysis, requires only a single realization to fit the non-stationary model. First we consider the building blocks of our model: discrete non-decimated 2D wavelets.

2.2 Discrete non-decimated 2D wavelets

We provide a brief description of wavelets here. The reader should consult Daubechies (1992), Vidakovic (1999) or Nason (2008) for further details. A set of wavelets is a set of functions $\{\psi_{j,k}(x)\}_{j,k \in \mathbb{Z}}$ that act as an (orthonormal) basis for functions f in a function space $L^2(\mathbb{R})$, say. The representation is given by $f(x) = \sum_{j,k} d_{j,k} \psi_{j,k}(x)$. If the basis functions are orthogonal then the coefficients can be obtained in the usual way, i.e.

$$d_{jk} = \int f(x) \psi_{jk}(x) dx. \quad (1)$$

The wavelets are all scalings and translations of a single function, called the *mother wavelet*, $\psi(x)$, defined by $\psi_{j,k}(x) = 2^{j/2} \psi(2^j x - k)$. The mother wavelet has several important properties: fast decay in time and frequency domain (often compactly supported in one domain) and zero integral. Hence, the wavelet coefficients $d_{j,k}$ of a function $f(x)$ convey information about that function at scale proportional to 2^j and location $2^{-j}k$.

Associated with a mother wavelet is a father wavelet, $\phi(x)$, which is similar to a kernel function such as that used in kernel density estimation. Whilst wavelet coefficients provide

information about the local oscillatory behaviour of a function the father coefficients store information about the multiscale mean behaviour of that function. The father wavelets satisfy a multiscale relation, called the *dilation equation*: $\phi(x) = \sum_k h_k \phi(2x - k)$, the wavelet $\psi(x)$ satisfies a similar equation with h_k replaced by g_k . Classes of mother/father wavelets can be characterised by a suitable choice of $\{h_k, g_k\}$. There are many families of wavelets. A particularly useful and famous set of compactly supported wavelets was developed by Daubechies (1988). We use this family extensively in this paper. A discrete wavelet transform (DWT) exists for sequence data: the pyramid algorithm due to Mallat (1989) carries out the DWT with $\mathcal{O}(n)$ computational effort and memory requirements.

Let $\{h_k, g_k\}$ be quadrature mirror filters associated with a particular Daubechies (1992) compactly supported continuous time wavelets. Let $j \in \mathbb{Z}^+$ be the scale (the negative of that in NvSK for a clearer presentation). Formulae (3) and (4) of NvSK introduced the discrete mother wavelets $\psi_j = (\psi_{j,0}, \dots, \psi_{j,L_j-1})$ where $L_j = (2^j - 1)(N_h - 1) + 1$ and N_h is the number of non-zero elements of $\{h_k\}$. We define the discrete *father* wavelets $\phi_j = (\phi_{j,0}, \dots, \phi_{j,L_j-1})$ in exactly the same way but replacing g_{n-2k} by h_{n-2k} of formula (3) of NvSK. As an example, the discrete Haar father wavelet filters at scales $j = 1, 2$ are $\phi_1 = (h_0, h_1) = (1/\sqrt{2})(1, 1)$ and $\phi_2 = (h_0^2, h_1 h_0, h_0 h_1, h_1^2) = \frac{1}{2}(1, 1, 1, 1)$. We now define the 2D discrete father and mother wavelets.

Definition 1. Let $\mathbf{k} = (k_1, k_2)$ where $k_1, k_2 \in \mathbb{Z}$. We define the **2D discrete wavelet filters**, $\{\psi_j^l\}$, as finite square matrices, of dimension, $L_j \times L_j$, as follows:

$$\psi_j^l = \begin{bmatrix} \psi_{j,(0,0)}^l & \cdots & \psi_{j,(0,L_j-1)}^l \\ \vdots & \vdots & \vdots \\ \psi_{j,(L_j-1,0)}^l & \cdots & \psi_{j,(L_j-1,L_j-1)}^l \end{bmatrix} \text{ for } l = h, v \text{ or } d,$$

where h, v and d denote the horizontal, vertical and diagonal directions; the elements are

$$\left. \begin{array}{l} \psi_{j,\mathbf{k}}^h = \phi_{j,k_1} \psi_{j,k_2} \\ \psi_{j,\mathbf{k}}^v = \psi_{j,k_1} \phi_{j,k_2} \\ \text{and } \psi_{j,\mathbf{k}}^d = \psi_{j,k_1} \psi_{j,k_2} \end{array} \right\} \text{ for } k_1, k_2 = 0, \dots, L_j - 1, \quad (2)$$

where $\psi_{j,k}, \phi_{j,k}$ are the 1D discrete wavelets. Similarly, 2D discrete father wavelets are defined by: $\phi_{j,\mathbf{k}} = \phi_{j,k_1} \phi_{j,k_2}$.

Example 1. For example, the discrete Haar wavelet in the diagonal decomposition direction at scales $j = 1, 2$ are given by:

$$\psi_1^d = \frac{1}{2} \begin{bmatrix} 1 & -1 \\ -1 & 1 \end{bmatrix} \text{ and } \psi_2^d = \frac{1}{4} \begin{bmatrix} 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \\ -1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 \end{bmatrix}.$$

As in 1D we can form the collection of *non-decimated* discrete wavelets by translations as follows: $\psi_{j,\mathbf{u}}^l(\mathbf{r}) = \psi_{j,\mathbf{u}-\mathbf{r}}^l$, for $j \in \mathbb{Z}^+$, directions l and *all* locations $\mathbf{u}, \mathbf{r} \in \mathbb{Z}^2$, in contrast to regular wavelets that are placed at dyadic locations. For further details see Nason and Silverman (1995), Nason (2008), or Unser (1995) for their use in texture analysis. We now introduce our wavelet model for random fields based on non-decimated wavelets.

2.3 Locally stationary wavelet random fields

We introduce a class of lattice processes composed of random mixtures of 2D discrete non-decimated wavelets. Our model is one particular possible 2D extension of the locally stationary wavelet (LSW) model of NvSK: the main structural difference between 1D and 2D cases is the introduction of the directional index l and lowest common scale J .

Definition 2. Let $\mathbf{R} = (R, S)$ where $R = 2^m, S = 2^n \geq 1$ for $m, n \in \mathbb{N}$ and set $J(R, S) \equiv \log_2\{\min(R, S)\}$ be the lowest common scale. Further, let $\mathbf{r} = (r, s)$ and $\mathbf{u} = (u, v)$ for $\mathbf{r}, \mathbf{u} \in \{0, \dots, R-1\} \times \{0, \dots, S-1\} = \mathcal{R}$. Then a class of *locally stationary 2D wavelet processes* (LS2W) is defined to be a sequence of stochastic processes defined on a regular grid and denoted by $\{X_{\mathbf{r};\mathbf{R}}\}_{\mathbf{r} \in \mathcal{R}}$ having the following representation in the mean-square sense:

$$X_{\mathbf{r};\mathbf{R}} = \sum_l \sum_{j=1}^{\infty} \sum_{\mathbf{u}} w_{j,\mathbf{u};\mathbf{R}}^l \psi_{j,\mathbf{u}}^l(\mathbf{r}) \xi_{j,\mathbf{u}}^l, \quad (3)$$

where the sum over l is over decomposition directions v, h and d . The decomposition consists of $\{w_{j,\mathbf{u};\mathbf{R}}^l\}$: amplitudes which quantify the contribution made to the process at location \mathbf{u} ; $\{\psi_{j,\mathbf{u}}^l(\mathbf{r})\}$: a collection of discrete non-decimated 2D wavelets and $\xi_{j,\mathbf{u}}^l$: a mean zero random orthonormal increment sequence satisfying

$$\mathbb{E} \left(\xi_{j,\mathbf{k}}^l \xi_{m,\mathbf{n}}^p \right) = \delta_{j,m} \delta_{\mathbf{k},\mathbf{n}} \delta_{l,p}.$$

Following Fryzlewicz (2003, Section 3.2.2), the process is constructed over *all* possible scales ($j = 1, \dots, \infty$) avoiding unnecessarily restrictive tail behaviours of key quantities introduced later. The LS2W model also obeys the following.

1. $\mathbb{E} \xi_{j,\mathbf{u}}^l = 0$, hence $\mathbb{E}(X_{\mathbf{r}}) = 0$. In real applications it is unlikely that a process will have a zero mean. To use our LS2W processes, and if a non-zero mean should exist, then it should be modelled, estimated and removed. There are a large number of ways in which mean removal could be accomplished, e.g., median polish, Cressie (1991), multivariate regression or wavelet shrinkage techniques, Vidakovic (1999).
2. For each $l \in h, v, d$ and scale $j \geq 1$ there exists a Lipschitz-continuous function (with respect to the L_1 norm) $W_j^l(\mathbf{z})$ where $\mathbf{z} \in (0, 1)^2$. These functions satisfy, $\forall j$ and l : (a) (*finiteness*) $\sum_l \sum_{j=1}^{\infty} |W_j^l(\mathbf{z})|^2 < \infty$, uniformly in $\mathbf{z} \in (0, 1)^2$; (b) (*stationarity control*) the Lipschitz constants L_j^l of W_j^l are uniformly bounded in j, l and

$$\sum_l \sum_{j=1}^{\infty} 2^{2j} L_j^l < \infty.$$

Note that this condition is subtly different to the 1D case; (c) (*linkage*) Let $\frac{\mathbf{u}}{\mathbf{R}} := (\frac{u}{R}, \frac{v}{S})$. Then there exists a sequence of constants C_j^l such that for each dimension set \mathbf{R} ,

$$\sup_{\mathbf{u}} \left| w_{j,\mathbf{u};\mathbf{R}}^l - W_j^l \left(\frac{\mathbf{u}}{\mathbf{R}} \right) \right| \leq \frac{C_j^l}{\max(R, S)}, \quad (4)$$

where for each $j = 1, \dots, J(\mathbf{R})$ the supremum is over all pairs of coordinates $\mathbf{u} \in \mathcal{R}$ and where $\{C_j^l\}$ fulfills $\sum_l \sum_{j=1}^{\infty} C_j^l < \infty$.

From now on we will drop the explicit dependence on \mathbf{R} although it is still assumed. In our 2D situation we track power in the covariance decomposition of $X_{\mathbf{r}}$ with respect to scale *and* direction. The smoothness assumptions on W_j^l control the variation of the $\{w_{j,\mathbf{u}}^l\}$ as a function of \mathbf{u} and hence the local stationarity of the process.

Example: Haar moving average (MA) fields

We use Haar wavelets to construct LS2W fields, and first define the generating fields.

Definition 3. Let $c \in \mathbb{R}$. A Haar MA field of order j_0 , in direction l_0 , is defined to be the LS2W process $X_{\mathbf{r}}^{j_0, l_0}$ generated by the Haar 2D non-decimated discrete wavelets with the following condition on the amplitudes:

$$w_{j,\mathbf{u}}^l = \begin{cases} c & \text{for } j = j_0, l = l_0, \\ 0 & \text{otherwise.} \end{cases}$$

For example, setting $c = \sigma$ in the definition for $j_0 = 1$ (finest scale) and $l_0 = d$ (diagonal direction), using Haar wavelets and setting the orthonormal increment sequence $\xi_{1,\mathbf{u}}^d = \epsilon_{\mathbf{u}}$ where $\{\epsilon_{\mathbf{u}}\}$ is a purely random process with mean zero and variance 1 gives

$$X_{\mathbf{r}}^{1,d} = \sigma \sum_{\mathbf{u}} \psi_{1,\mathbf{u}-\mathbf{r}}^d \xi_{1,\mathbf{u}}^d = \sigma (\epsilon_{r,s} - \epsilon_{r,s+1} - \epsilon_{r+1,s} + \epsilon_{r+1,s+1})/2, \quad (5)$$

where $\mathbf{r} = (r, s)$. Figure 1(a) shows a realization of the $X_{\mathbf{r}}^{1,d}$ Haar MA field: it shows fine ‘‘diagonal’’ detail as it is built from the finest scale wavelets in the diagonal direction. Figure 1(b) shows coarser detail from scale $j_0 = 2$ horizontal ($l_0 = h$) and vertical wavelets ($l_0 = v$): a realization from the addition of two Haar MA fields with $w_{2,\mathbf{u}}^h = w_{2,\mathbf{u}}^v = \sigma$ (with all other w zero) giving

$$\begin{aligned} X_{\mathbf{r}}^2 &= \sigma \sum_{\mathbf{u}} \psi_{2,\mathbf{u}}^h(\mathbf{r}) \xi_{2,\mathbf{u}}^h + \sigma \sum_{\mathbf{u}} \psi_{2,\mathbf{u}}^v(\mathbf{r}) \xi_{2,\mathbf{u}}^v \\ &= \frac{\sigma}{2} \{ (\epsilon_{r,s} + \epsilon_{r,s+1} + \epsilon_{r+1,s} + \epsilon_{r+1,s+1}) - (\epsilon_{r+2,s+2} + \epsilon_{r+2,s+3} + \epsilon_{r+3,s+2} + \epsilon_{r+3,s+3}) \} \end{aligned} \quad (6)$$

Haar MA fields are special cases of the MA fields due to Haining (1978), see Moore (1988) or Cressie (1991) for further details. Any 2D MA field can be represented as the linear combination of Haar MA fields, but not uniquely as the non-decimated representation is



Figure 1: 2D Haar MA fields. *Left*: $j = 1$, d detail; *Right*: $j = 2$, h & v detail

over-determined. The representation follows since the (decimated) Haar wavelets form a basis for the $l^2(\mathbb{Z}^2)$ sequence space. More generally, Daubechies MA fields can be built using Daubechies' (1992) wavelets. Further, because we are using shift-equivariant non-decimated wavelets, LS2W includes a large class of correlated processes. In particular, all stationary processes satisfying $\sum_{\tau} |c(\tau)| < \infty$.

Whilst the examples given above are stationary processes, the real potential of the LS2W methodology lies in its ability to capture locally stationary behaviour. Figure 2 shows a realization obtained by juxtaposing four stationary LS2W processes: $X_{\mathbf{r}} = \sum_{j=1}^4 \sum_{\mathbf{u}} w_{j,\mathbf{u}}^d \psi_{j,\mathbf{u}}^d(\mathbf{r}) \xi_{j,\mathbf{u}}^d$, where $\{\psi_{j,\mathbf{u}}^l\}$ is the set of 2D Haar non-decimated discrete wavelets, and fixing

$$w_{j,[2^j \mathbf{z}]}^d = \begin{cases} \sigma & \text{for } j = 1, \mathbf{z} \in (0, 1/2) \times (0, 1/2); \\ \sigma & \text{for } j = 2, \mathbf{z} \in (1/2, 1) \times (0, 1/2); \\ \sigma & \text{for } j = 3, \mathbf{z} \in (0, 1/2) \times (1/2, 1); \\ \sigma & \text{for } j = 4, \mathbf{z} \in (1/2, 1) \times (1/2, 1); \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

Strictly, this example does not meet our LS2W conditions above as it contains jump discontinuities. However, the definition could be extended along the lines of the 1D extension proposed by Van Bellegem and von Sachs (2008) which does permit such jumps.

The next section introduces local autocovariance of LS2W processes.

2.4 Local wavelet spectra

The covariance structure of the LS2W process in Figure 2 clearly varies from one quadrant to the next. The following quantity measures the local power of an LS2W process at a specific (rescaled) location $\mathbf{z} \in (0, 1)^2$, scale j and direction l .

Definition 4. Let $X_{\mathbf{r}}$ be LS2W. The *local wavelet spectrum* (LWS) of $X_{\mathbf{r}}$ is defined by

$$S_j^l(\mathbf{z}) = |W_j^l(\mathbf{z})|^2 \quad (8)$$

for $\mathbf{z} \in (0, 1)^2$, $j \in 1, \dots, J$, and $l \in \{h, v \text{ or } d\}$.




Figure 2: Realization of the nonstationary LS2W process given in (7), $R = S = 2^9$. The texture detail ranges from finest (bottom left) to coarsest (top right).

The LWS is a 2D extension of the evolutionary wavelet spectrum from NvSK and an analogue of the stationary stochastic process spectrum.

Example (continued): Haar MA fields (concatenated nonstationary)

The nonstationary LS2W example above has LWS equal to equation (7) with $w_{j,2^j\mathbf{z}}^d$ replaced by $S_j^d(\mathbf{z})$ and σ by σ^2 . Eckley (2001, Example 3.6) gives simulation results for spectral estimation with this process.

2.5 Covariance of LS2W processes

Since LS2W processes are built from wavelets it follows that their covariance can be represented in terms of the covariance functions of 2D discrete non-decimated wavelets. We define the wavelet covariance functions, and then LS2W local covariance.

Definition 5. Let $j \in \mathbb{N}$, $l \in \{v, h, d\}$ and $\boldsymbol{\tau} \in \mathbb{Z}^2$. Then the **autocorrelation (ac) wavelet** of a 2D discrete wavelet family $\{\psi_{j,\mathbf{k}}^l\}$ is given by

$$\Psi_j^l(\boldsymbol{\tau}) = \sum_{\mathbf{v} \in \mathbb{Z}^2} \psi_{j,\mathbf{v}}^l(\mathbf{0}) \psi_{j,\mathbf{v}}^l(\boldsymbol{\tau}). \quad (9)$$

The 2D ac wavelets are separable because the discrete wavelets are from (2), i.e. in the horizontal, vertical and diagonal directions:

$$\Psi_j^h(\boldsymbol{\tau}) = \Phi_j(\tau_1)\Psi_j(\tau_2), \quad \Psi_j^v(\boldsymbol{\tau}) = \Psi_j(\tau_1)\Phi_j(\tau_2), \quad \Psi_j^d(\boldsymbol{\tau}) = \Psi_j(\tau_1)\Psi_j(\tau_2), \quad (10)$$

where $\boldsymbol{\tau} = (\tau_1, \tau_2)$, and Ψ_j, Φ_j are the 1D discrete ac wavelet and father wavelets from NvSK. The 2D discrete ac father wavelet is similarly given by $\Phi_j(\boldsymbol{\tau}) = \Phi_j(\tau_1)\Phi_j(\tau_2)$. Refer to Eckley and Nason (2005) for further details on a.c. wavelets.

Example: 2D Haar ac wavelets

The 1D Haar ac wavelet (see NvSK) is $\Psi_H(u) = 1 - 3|u|$ for $|u| \in [0, 1/2]$, and $|u| - 1$ for $|u| \in (1/2, 1]$. The 2D Haar ac wavelet is given by $\Psi_j^l(\tau_x, \tau_y) =$

$\Psi_H^l(2^{-j}|\tau_x|, 2^{-j}|\tau_y|)$, where $\Psi_H^l(\mathbf{u})$ is constructed from the separability relations in (10) depending on l . Figure 3 shows $\Psi_5^l(\tau_1, \tau_2)$ for the Haar family.




Figure 3: 2D Haar autocorrelation wavelet at scale 5.

For stationary processes it is well-known that the autocovariance is the Fourier transform of the spectrum. Is there a relationship between the covariance of a LS2W process and its local wavelet spectra? Yes! The covariance of a LS2W process tends asymptotically to the “wavelet transform” of the LWS, $C(\mathbf{z}, \boldsymbol{\tau})$, given by the following definition.

Definition 6. Define the **local covariance** (LCV), $C(\mathbf{z}, \boldsymbol{\tau})$, of a given LS2W process with LWS $\{S_j^l(\mathbf{z})\}$, to be

$$C(\mathbf{z}, \boldsymbol{\tau}) = \sum_l \sum_{j=1}^{\infty} S_j^l(\mathbf{z}) \Psi_j^l(\boldsymbol{\tau}), \quad (11)$$

where $\boldsymbol{\tau} \in \mathbb{Z}^2$ and $\mathbf{z} \in (0, 1)^2$.

Let $C_{\mathbf{R}}(\mathbf{z}, \boldsymbol{\tau})$ be the autocovariance of a LS2W process $X_{\mathbf{r}}$, i.e. $C_{\mathbf{R}}(\mathbf{z}, \boldsymbol{\tau}) = \text{Cov}(X_{[\mathbf{z}\mathbf{R}]}, X_{[\mathbf{z}\mathbf{R}] + \boldsymbol{\tau}})$. The following proposition shows that $C_{\mathbf{R}}$ asymptotically converges to C in (11).

Proposition 1. Let $C_{\mathbf{R}}$ be the autocovariance of a LS2W process $X_{\mathbf{r}}$ and C as in Definition 6. Then $|C_{\mathbf{R}}(\mathbf{z}, \boldsymbol{\tau}) - C(\mathbf{z}, \boldsymbol{\tau})| = O\{\min(R, S)^{-1}\}$ as $R, S \rightarrow \infty$, uniformly in $\boldsymbol{\tau} \in \mathbb{Z}^2$ and $\mathbf{z} \in (0, 1)^2$.

Proof: See Eckley, Nason and Treloar (2008).

If $X_{\mathbf{r}}$ is stationary then $S_j^l(\mathbf{z})$ is constant over \mathbf{z} and C is the usual autocovariance function. This is illustrated by the following example.

Example (continued): Haar MA fields

The (stationary) process, $X_{\mathbf{r}}^{1,d}$ given in (5) has autocovariance:

$$C_{X^1}(\tau_1, \tau_2) = \sigma^2 \{ \delta_{\tau_1,0} \delta_{\tau_2,0} - (\delta_{\tau_1,0} \delta_{\tau_2,1} + \delta_{\tau_1,1} \delta_{\tau_2,0} + \delta_{\tau_1,0} \delta_{\tau_2,-1} + \delta_{\tau_1,-1} \delta_{\tau_2,0})/2 + (\delta_{\tau_1,1} \delta_{\tau_2,1} + \delta_{\tau_1,-1} \delta_{\tau_2,-1} + \delta_{\tau_1,-1} \delta_{\tau_2,1} + \delta_{\tau_1,1} \delta_{\tau_2,-1})/4 \}.$$

In other words, X^1 has a sparse covariance representation in terms of equation (11): $C_{X^1}(\tau_1, \tau_2) = \sigma^2 \Psi_1^d(\tau_1, \tau_2)$, which does not depend on \mathbf{z} and is exact (not asymptotic).

2.6 Uniqueness of the covariance representation

As in NvSK the question of whether (11) is invertible arises: can we represent the spectrum in terms of the LCV (well-known for stationary processes)? This hinges on the invertibility of the inner product matrix of the autocorrelation wavelets, A_J . To demonstrate this we first define the inner product matrix of 2D discrete autocorrelation wavelets. This matrix differs from the 1D case as the 2D wavelets have a directional component. To simplify notation, we introduce a new multi-index combining scale and direction.

Definition 7. A 2D wavelet (or autocorrelation wavelet) at scale j and direction l may be indexed by $\eta(j, l) = j + g(l)$, where $g(l) = 0, J, 2J$ for $l = v, h, d$ for all $j = 1, \dots, J$.

Hence the first J entries of η correspond to vertical wavelets, the next J to horizontal and the last J to diagonal. To simplify this we will omit the dependency on j and l . We now define the inner product matrix of discrete ac wavelets as follows.

Definition 8. Define the operator $A = (A_{\eta,\nu})_{\eta,\nu \geq 1}$ by

$$A_{\eta,\nu} = \langle \Psi_{\eta}, \Psi_{\nu} \rangle = \sum_{\tau} \Psi_{\eta}(\tau) \Psi_{\nu}(\tau). \tag{12}$$

Define the $3J$ -dimensional matrix $A_J = (A_{\eta,\nu})_{\eta,\nu=1,\dots,3J}$, where $J = \log_2(\min(R, S))$.

The following theorem demonstrates the invertibility of A_J .

Theorem 1. For any compactly supported Daubechies wavelet, the family of discrete 2D autocorrelation wavelets $\{\Psi_{\eta}\}$ is linearly independent. Hence,

1. the operator A is invertible (since all of its eigenvalues are positive) and for each $J \in \mathbb{N}$, the norm $\|A_J^{-1}\|$ is bounded above.
2. the LWS is uniquely defined given the corresponding LS2W process.

Proof: See Eckley, Nason and Treloar (2008).

Invertibility permits us to show that the spectrum can be represented in terms of the LCV:

Corollary 1. *The inverse formula of (11) is*

$$S_j^l(\mathbf{z}) = \sum_{\eta_1} A_{\eta_1}^{-1} \sum_{\boldsymbol{\tau}} C(\mathbf{z}, \boldsymbol{\tau}) \Psi_{\eta_1}(\boldsymbol{\tau}). \quad (13)$$

Proof: See Eckley, Nason and Treloar (2008).

Theorem 2 in NvSK conjectured that the smallest eigenvalue of their infinite order operator A was bounded away from zero (proved for Haar and Shannon wavelets). This more stringent property is needed for two important results required for the estimation of the LWS (Corollary 1 and Proposition 3 of NvSK). Whilst our infinite order operator A in (12) differs in structure to that considered by NvSK, we conjecture that its smallest eigenvalue is also bounded away from zero. The result is easy to show for Shannon wavelets.

3 Estimating the LWS

Having found a measure which provides a local direction-scale decomposition of power, it is natural to enquire how one can estimate this quantity, given the prior specification of the underlying wavelet family. The issue of what happens when one uses an alternative wavelet family to that which underlies the process is left as an avenue for future work.

Recall from stationary theory that an estimate of the spectral density function is given by the squared absolute value of the Fourier transform. As in NvSK, the estimator which we propose for the LWS is founded upon the collection of squared empirical wavelet coefficients – the local wavelet periodogram.

Definition 9. Let $\{X_r\}$ be a LS2W process as defined in Definition 2. The empirical wavelet coefficients of the process are given by $d_{j,\mathbf{u}}^l \equiv \sum_{\mathbf{r}} X_{\mathbf{r}} \psi_{j,\mathbf{u}}^l(\mathbf{r})$.

We are now in a position to define the local wavelet periodogram (LWP).

Definition 10. The LWP of a LS2W process $\{X_{\mathbf{r}}\}$ is defined as

$$I_{j,\mathbf{u}}^l \equiv |d_{j,\mathbf{u}}^l|^2. \quad (14)$$

As we demonstrate in Theorem 2, the LWP is a biased estimator of the LWS. However the form of this bias suggests a transformation of the spectra which produces an asymptotically unbiased estimate of the LWS.

Theorem 2. Let $\mathbf{z} = (z_1, z_2)$, $\mathbf{R} = (R, S)$ and $[\mathbf{z}\mathbf{R}] = ([\mathbf{z}_1 R], [\mathbf{z}_2 S])$ where $R = 2^J, S = 2^K$ for some $J, K \in \mathbb{N}$. Further, assume that the $\{\xi_{\eta,\mathbf{r}}\}$ are Gaussian. Then,

$$\mathbb{E}(I_{\eta,[\mathbf{z}\mathbf{R}]}) = \sum_{\eta_1} A_{\eta_1} S_{\eta_1}(\mathbf{z}) + O\left(\frac{1}{\min\{R, S\}}\right). \quad (15)$$

Proof: See Eckley, Nason and Treloar (2008).

Thus the LWP estimate of the LWS at a given (j, l) -pair is a weighted sum of the LWS at all its locations. An example demonstrating this bias can be seen in Table 1. Note how in the case of the third finest scale LWP in the vertical decomposition direction, $I_{3, [\mathbf{zR}]}^v$, the estimator is a mix of contributions from various directions and scales. In particular, power leaks from $(3, v)$ across into the diagonal decomposition direction.

Clearly, without correction, the redundancy of the non-decimated wavelet transform (NDWT) induces a spread of power into other directions and scales. However, if we denote the vector of periodograms, $\mathbf{I}(\mathbf{z}) = \{I_{\eta, [\mathbf{zR}]}\}$, and define the vector of corrected LWPs to be given by $\mathbf{L}(\mathbf{z}) = A_J^{-1}\mathbf{I}(\mathbf{z})$, then we obtain an asymptotically unbiased estimator of the LWS:

$$\mathbb{E}(\mathbf{L}(\mathbf{z})) = \mathbf{S}(\mathbf{z}) + O\left(\frac{1}{\min\{R, S\}}\right) \quad (16)$$

Table 1: Bias weights, A_{η_1} , (to 3 d.p.) entering into the LWP estimate of $S_3^v(z)$. For an unbiased estimate only scale 3 vertical should be non-zero.

Direction	Scale 1	Scale 2	Scale 3	Scale 4	Scale 5
Vertical	0.703	3.797	15.453	13.793	7.573
Horizontal	0.203	0.797	1.891	2.793	2.073
Diagonal	0.047	0.422	3.953	8.379	6.220

The following definition will prove useful when considering the covariance structure of the wavelet periodogram.

Definition 11. Define

$$\alpha_{j_1, j_2}^{l_1, l_2}(\mathbf{u}_1, \mathbf{u}_2) = \sum_{\mathbf{r}} \psi_{j_1, \mathbf{u}_1}^{l_1}(\mathbf{r}) \psi_{j_2, \mathbf{u}_2}^{l_2}(\mathbf{r}). \quad (17)$$

In effect, this is a form of ‘‘cross-correlation’’ between two wavelets of the same family at (possibly) different scales and directions, centred on different locations. Using this identity, we can explore the covariance structure of the (uncorrected) LWP.

Theorem 3. Assume that the $\{\xi_{\eta, \mathbf{r}}\}$ are again Gaussian. Then the covariance between $I_{j_1, \mathbf{p}}^{l_1}$ and $I_{j_2, \mathbf{q}}^{l_2}$ may be expressed as follows:

$$\text{Cov}(I_{j_1, \mathbf{p}}^{l_1} I_{j_2, \mathbf{q}}^{l_2}) = 2 \left\{ \sum_{l_0} \sum_{j_0} \sum_{\mathbf{u}_0} (w_{j_0, \mathbf{u}_0}^{l_0})^2 \alpha_{j_1, j_0}^{l_1, l_0}(\mathbf{p}, \mathbf{u}_0) \alpha_{j_2, j_0}^{l_2, l_0}(\mathbf{q}, \mathbf{u}_0) \right\}^2.$$

Thus the correlation between these quantities decreases with increasing distance between location \mathbf{p} at scale-direction (j_1, l_1) and the location \mathbf{q} at (j_2, l_2) . In particular, when $j_1 = j_2$, the covariance is zero when $\|\mathbf{p} - \mathbf{q}\|$ exceeds the overlap of the corresponding

wavelets support. Moreover

$$\begin{aligned}\text{Var}(I_{j,\mathbf{p}}^l) &= 2\mathbb{E}(I_{j,\mathbf{p}}^l)^2 \\ &= 2\left(\sum_{\eta_1} A_{\eta_1} S_{\eta_1}([\mathbf{p}/\mathbf{R}])\right)^2 + O\left(\frac{2^{j(\eta)}}{\min(R, S)}\right),\end{aligned}\quad (18)$$

where $j(\eta) \equiv \eta - \lfloor \frac{\eta-1}{J} \rfloor J$ simply denotes the scale element of $\eta(j, l)$.

Proof: See Eckley, Nason and Treloar (2008).

The above demonstrates that the uncorrected LWP has asymptotically non-vanishing variance. Hence, by construction, the *corrected* LWP will also have an asymptotically non-vanishing variance, thus paralleling the traditional stationary case. Consequently, our estimates of the LWS will be smoothed to obtain consistency.

Several smoothing approaches could be used in this instance, for example kernel smoothing or a moving average approach. However as images are characterised by edges it would appear prudent to use a smoothing scheme which has the ability to deal efficiently with such features. Assuming that the innovations $\{\xi_{j,\mathbf{u}}^l\}$ are Gaussian it follows that, upon squaring, each element of the wavelet periodogram has a χ^2 -distribution. Correcting, to obtain an asymptotically unbiased estimate of the LWS (as suggested by Theorem 3) leads to a complex correlated distribution for the LWP. Thus, we follow NvSK and suggest firstly performing wavelet shrinkage of the χ^2 -distributed periodogram prior to correction by A^{-1} . A detailed description of how one may smooth using an orthonormal second-stage wavelet basis $\tilde{\psi}_{l,m}$ is provided by von Sachs, Nason and Kroisandt (1997). Briefly, smoothing is performed by implementing a non-linear thresholding of the raw (uncorrected) periodogram, $I_\eta(\mathbf{z})$, and then inverting the smoothed transformation to obtain the estimate $\tilde{L}_\eta(\mathbf{z})$. This is the approach adopted within the LS2W software package developed by Eckley and Nason (2009), who also provide details on various approaches which can be used to visualise the LWP collection.

4 Texture description and analysis

We now consider the application of the LS2W modelling approach to texture analysis. In laymen's terms, texture is the visual character of an image region whose structure is, in some sense, regular: for example the appearance of a woven material.

Texture frequently possesses structure on many different scales. Thus, when modelling the structure of a textured image, an attempt should be made to incorporate this multiscale reality. A model, such as that afforded by the LS2W approach, which provides a multiscale decomposition of the covariance structure of a textured image would therefore appear desirable. A brief introduction to statistical texture analysis is provided below, with particular emphasis on wavelet-based approaches which have recently appeared in the literature. For a more comprehensive review, the reader is referred to Petrou (2006).

4.1 Statistical approaches to texture description

Comprehensive reviews of the statistical approach to texture analysis are provided by Haralick (1979; 1986), Tuceryan and Jain (1999), Tomita and Tsuji (1990, Chapter 2) and Petrou (2006). Reed and du Buf (1992) review feature extraction techniques for unsupervised applications whilst Randen and Husøy (1999) provide a comparative review of various filtering-based approaches to feature extraction.

Perhaps the most familiar statistical techniques are those based upon the autocorrelation function and spectrum. Typically, statistics including average values of energy within ring or wedge functions of frequency are considered. These provide features relating to coarseness and directionality respectively (see Weszka (1976) for further details). An alternative approach is to consider texture in terms of edgeness per unit area, see Davies and Mitchie (1980).

Haralick *et al.* (1973) present a general procedure for extracting textural properties based upon the co-occurrence matrix of an image. This matrix, $P_{d,\phi}(a,b)$, measures the number of occurrences with which two pixels, of gray levels a and b respectively, appear in \mathbf{R} separated by a distance d in direction ϕ . Various measures such as energy, entropy, contrast and correlation may be derived from the co-occurrence matrices, these features subsequently being used for texture classification etc. Discrete sine, cosine and Hadamard transforms are all examples of a local linear transform (LLT). With several potential LLTs available for any given problem, Unser (1986) considers the issue of transform selection for a given application.

4.2 Multiscale approaches to texture analysis

Recent psycho-visual research has indicated that the human and mammalian visual systems process images in a multiscale manner, preserving both local and global information; see Daugman (1990), Reed and Wechsler (1990) or Field (1999) for example. Such findings have provided a strong motivation for the development of texture analysis techniques founded upon multiscale methods.

Initial multiscale approaches to texture analysis were based upon Gabor functions, see for example Turner (1986), Bovic, Clark and Geisler (1990) and Dunn and Higgins (1995). However Unser (1995) provides compelling arguments *against* such an approach, highlighting potential disadvantages including computational intensity

The use of the discrete wavelet transform (DWT) for texture analysis was first suggested by Mallat (1989). This transform is appealing as it is well localised and permits a decomposition into three different directions: vertical, horizontal and diagonal. However, as Chang and Kuo (1993) reason, a potential disadvantage of using the DWT for texture analysis is that it focuses on the progressive analysis of the *low*-frequency smooths. Thus, the DWT does not always provide a suitably refined partition of the middle frequencies. To combat this, Chang and Kuo (1993) suggest the use of the “tree structured” or wavelet packet transform. Similar ideas are proposed by Saito and Coifman (1995) and Laine and Fan (1993).

Whilst appealing, methods such as the DWT and the Discrete Wavelet Packet Transform lack translational equivariance (TE). Put simply, the consequence of non-TE is that a simple integer shift of the input signal frequently results in a non-trivial change in the DWT of the signal. This is clearly undesirable. To remedy this issue, Unser (1995) proposes the use of the Discrete Wavelet Frame (DWF), a form of non-decimated wavelet transform, for texture analysis. Van de Wouwer, Scheunders and Van Dyck (1999) consider the application of the discrete (undecimated) wavelet transform to texture analysis, introducing two new feature sets: (i) based on parameter estimates for a Weibull distribution of the wavelet detail coefficients; (ii) motivated by the work of Haralick *et al.* (1973), calculating co-occurrence matrices of the wavelet detail images.

To overcome problems of translation invariance and poor directional selectivity within the DWT, novel multiscale transforms such as the non-decimated wavelet packet transform, dual-tree complex wavelet transform (Kingsbury (1999)) and the steerable pyramid (Simoncelli (1995)) have been used for various texture analysis tasks. The work by Portilla and Simoncelli (2000) is particularly interesting, resulting in excellent texture analysis and synthesis performance. Their steerable pyramid is rotationally and translation invariant, and like the LS2W model, is based on an overcomplete system. We prefer a more classical statistical-based approach which specifies a model, and then develops an unbiased estimator for that model, whereas the texture analysis/synthesis work by Portilla and Simoncelli (2000) appears to model aspects of the overcomplete coefficients probability structure and then use a method to synthesize textures that agree with the estimated probability structure. In our model language this would be equivalent to working with the I periodogram rather than the S spectrum. We prefer the latter because the model specification is in terms of S .

Research by Baraniuk and collaborators have focused upon hidden Markov tree modelling of the structure contained within wavelet transforms: for example Crouse, Nowak and Baraniuk (1998), Romberg, Choi and Baraniuk (2001), Venkatachalam, Choi and Baraniuk (2000). Such models can capture the key features of many real world images, for example the persistent nature of discontinuities in the wavelet domain. However, the application of such approaches can be computationally expensive, see Romberg *et al.* (2001). To combat such expense, it is often convenient to reduce the number of model parameters, assuming that within any given scale, the parameters are constant over location.

Remark 1. The measures used by Unser (1995) for texture classification are similar to those which we will consider in Section 5 when we apply the LS2W model to various texture analysis problems. Both sets of measures are based on translation equivariant wavelet transforms. However, rather than being motivated by a measure of energy or entropy, our measure is model-based and our modelling framework permits us to recognize a statistical bias and that, as a consequence, power in scales leaks across to other scales and directions. Hence we can correct for this bias to obtain superior results in applications and attach meaningful interpretations to spectral quantities. Finally, our model permits us to synthesize texture in a controlled, model-based way. However our texture synthesis is not as general, as that described by Portilla and Simoncelli (2000) who make use of a steerable wavelet pyramid which has the benefit of rotational invariance (and shares TE with our method).

5 The LS2W model and texture analysis

The LS2W model developed in Section 2 provides a rigorous stochastic framework upon which we can build a texture discrimination/classification scheme.

Given a textured image, T_I , of dimension $2^J \times 2^J$, the collection of (smoothed, corrected) local wavelet periodograms, $\{\tilde{\mathbf{L}}(\mathbf{z})\}$, forms an array of dimension $3J \times 2^J \times 2^J$. As a first step to investigating the potential of the LS2W approach to texture analysis, we consider the following statistic, one of many which could be based upon this measure:

$$\mathbf{t}(T_I) = \sum_{\mathbf{z}} \tilde{\mathbf{L}}(\mathbf{z}) = \sum_{\mathbf{z}} A_J^{-1} \tilde{\mathbf{I}}([\mathbf{z}\mathbf{R}]), \quad (19)$$

where $\tilde{\mathbf{I}}$ denotes the smoothed (uncorrected) local wavelet periodogram. Any given element, $\{\mathbf{t}(T_I)_{\eta}\}_{\eta(j,l)}$, provides a measure of the contribution made to the overall local variance structure at scale j within direction l . This measure is similar to the ‘‘channel-variance’’ proposed by Unser (1995). However, whilst Unser’s feature set is motivated by the conservation of energy within a tight wavelet frame, no consideration is made of how the redundancy of the DWF can affect estimates of local spectral features.

In the remainder of this section, we consider the application of the LS2W approach to two specific texture problems. We begin by focusing on its potential to discriminate and classify between subtly different textures encountered in pilled textile images. The second problem focuses on the more complicated issue of non-stationary texture classification. We will compare our results with those of other suitable approaches. Whilst the Brodatz collection has become a standard in the texture analysis literature, we do not report any results based on this collection here. Full results of our tests on this set are reported in Eckley (2001).

5.1 Exploratory analysis of pilled material images

Before we begin we should warn the reader that since the images are approximately stationary the Fourier techniques beat the wavelet methods, although with an effective classification scheme wavelets do nearly as well. This is important since we want wavelets to do well in the stationary case but Fourier is still the optimal paradigm in this setting. The true power of our LS2W methodology is shown with the non-stationary problem in the next section.

Six samples of identical material were buffed to varying degrees in an attempt to simulate different levels of garment wear. The effect of this buffing is to induce pilling, a building up of fibrous balls on the surface of the material. As can be seen in Figure4, certain materials have a very fine level of pilling (for example Figure4 (1)) whilst others are heavily pilled (Figure4 (6)).

Interestingly, some of these samples are very difficult to discriminate between visually. To investigate the ability of the LS2W approach to discriminate between these different textures, fifty sub-images of dimension 128×128 were randomly sampled from the left hand half of each image. For each sub-image, feature sets based upon




Figure 4: Images of materials pilled to varying degrees. Image (1) contains a fine pill material, whilst image (6) contains heavy pilling. Images provided by Unilever Research.

Method I: using Daubechies' Extremal Phase wavelets ($N=3$) and smoothed using Daubechies Least Asymmetric ($N=6$) wavelet;

Method II: the uncorrected non-decimated wavelet transform;

Method III: the discrete wavelet transform;

Method IV: Fourier rings (of 10 frequency units) were evaluated.

The linear partition of the Fourier frequency space used was thought to be reasonable for this initial study, being neither particularly fine nor coarse. Other choices of partition could consist of a fine linear partition of the space or a logarithmic partition, thus mimicking the division performed by wavelets. Daubechies Extremal Phase ($N=3$) wavelets were used for all wavelet-based measures.

Figure 5 displays a plot of the first two linear discriminant axes for the LS2W feature set. Note how the different pill levels span the plane: heaviest pill on the left and lightest pills on the right. The different classes are reasonably well separated, the analysis even being able to separate pill levels 5 and 6, two images which appear very similar to the eye. However, it should be noted that pill levels 3 and 4 overlap. To view discriminant plots associated with uncorrected NDWT, DWT and Fourier based features, refer to Eckley (2001, p. 103).

With such subtle differences between the images displayed in Figure 4, it is interesting to see whether the various feature extraction schemes can provide measures which permit reasonable classification rates. To this end, a test set of fifty sub-images of dimension




Figure 5: Linear discriminant analysis plots for measures of the Pill images based on the LS2W model.

128×128 were randomly sampled from the right half of each pill image with Method I - IV features being generated for each sub-image. These feature vectors were subsequently used to classify the sub-images to a pill class.

To begin with, a single stage classification scheme based upon the first four linear discriminant variables was considered. Each sub-image was classified using a minimum (Euclidean) distance rule. In the event that the distances between a sub-image and two (or more) texture classes were equal, the sub-image was deemed to be unclassified. The results of this approach are displayed in Table 2. As can be seen, barely half the sub-images are classified correctly by the multiscale methods — the LS2W approach achieving the best results of the three. Note however, that the Fourier approach classifies approximately two thirds of the sub-images correctly. These comparatively poor misclassification rates are not particularly surprising, for texture classes 2, 3 and 4 are poorly separated by linear discriminant analysis.

Noticing that it is difficult to discriminate between, for example, pill levels 3 and 4 in the linear discriminant analysis plots associated with the multiscale approaches, it is natural to consider a two-stage scheme in an attempt to improve classification performance. Such an approach yields improved classification results. It is perhaps not unsurprising that the results for the various methods are similar, for these images have a regular form. Consequently their spectral properties in the wavelet domain will also be regular, implying that the underlying process is in some sense stationary. We would therefore expect these textures to be well-discriminated by Fourier features. See Eckley (2001) for further details.

Table 2: % of Pill textures classified correctly with one- and two-step classification algorithms.

Method	% Correct (one step)	% Correct (two step)
I	57.7	70.7
II	51.7	66
III	54	65.7
IV	66	72.3

5.2 Non-stationary texture classification

The power of the LS2W modelling approach lies in its ability to analyse images whose covariance structure is *locally* stationary. In other words, it is well suited to the analysis of images whose covariance structure is *globally* non-stationary, but *stationary* within a local region. Crucially, the LS2W approach is able to correct artefacts which arise as a consequence of the inherent redundancy of the NDWT, the transform used in the estimation of the spectral structure of an image. The result of this correction is that we are able to stop power spreading across scales and directions. This is in stark contrast to using the uncorrected techniques.

Suppose a certain tile making process generates two texture types, T1 and T2 (see Figure 6). T1 represents a desirable tile type whilst T2 is deemed to be a spoiled tile. The task therefore is to find an approach which is able to achieve a high rate of correct classification.




Figure 6: Simulated examples of non-stationary textures.

Two classification approaches are considered, the first being based upon our LS2W model whilst the second uses the NDWT. The LS2W classification approach is structured as follows: For each of 25 realisations of tiles T1 and T2, calculate the local wavelet periodogram using the Haar transform, smoothing each periodogram using Daubechies Extremal Phase (N=4) wavelets. Then calculate the mean local wavelet periodogram

structure within each tile type, thus obtaining two spectral models, $\tilde{\mathbf{I}}_{T1}$ and $\tilde{\mathbf{I}}_{T2}$, of each tile's local wavelet spectral structure

A further 50 realisations of each tile type, $\{T_i\}_{i=1,\dots,100}$, are then used as a test set for classification purposes. For each test case, calculate the LWP, again using the squared detail coefficients of the Haar NDWT smoothed using the Daubechies Extremal Phase (N=4) wavelets. A tile T_i is then assigned to type T1 if

$$\sum_{j,l,\mathbf{u}} \left(\tilde{L}_{j,l,\mathbf{u};T1} - \tilde{L}_{j,l,\mathbf{u};T_i} \right)^2 < \sum_{j,l,\mathbf{u}} \left(\tilde{L}_{j,l,\mathbf{u};T2} - \tilde{L}_{j,l,\mathbf{u};T_i} \right)^2$$

and type T2 if

$$\sum_{j,l,\mathbf{u}} \left(\tilde{L}_{j,l,\mathbf{u};T1} - \tilde{L}_{j,l,\mathbf{u};T_i} \right)^2 > \sum_{j,l,\mathbf{u}} \left(\tilde{L}_{j,l,\mathbf{u};T2} - \tilde{L}_{j,l,\mathbf{u};T_i} \right)^2.$$

An equivalent approach is adopted using the squared detail coefficients of an unsmoothed, Haar non-decimated wavelet transform of the realisations.

Recall that in Section 5.1 the LS2W and (uncorrected) NDWT approaches yielded similar classification rates. This was due to the original images being stationary. In this case, the *uncorrected* NDWT method is only able to classify 62% of tiles correctly, but the corrected version classifies all correctly. The reason for this is that the inherent redundancy of the NDWT causes power to leak across directions and into lower scales, thus making discrimination between the two tile types on the basis of their detail coefficients difficult. The LS2W approach corrects for this leakage and therefore attains a higher classification rate.

A further issue to consider here is that the wavelet used for generating the textures and analysing the textures is the same: Haar. In practice, of course, the generating wavelet might be different or, more realistically, not anything necessarily to do with wavelets. However, even in this artificial simulation situation, the matching of the wavelets is not a significant issue. The reason being that the model is comprised of wavelets, but with *random* coefficients. Hence, there is no reason that resulting process should 'look' anything like the underlying wavelet (e.g. a Brownian motion does not look particularly 'boxy'). This is unlike the more familiar situation of nonparametric regression using wavelets where the model is an additive 'signal+noise' model and where, if the noise level is small, the noisy function 'looks like' a collection of wavelets, and hence it can be critical which wavelet is chosen to analyze the signal.

6 Conclusions and further work

Wavelet methods have been applied to many branches of statistics, from density estimation to time series analysis. In a departure from these comparatively established areas of research, this paper has considered the application of wavelets to the modelling of locally stationary random fields which lie on a regular grid. We introduced the LS2W model, which permits a local decomposition of the covariance structure into various scale contributions

within certain directions. A wavelet analogue of the Fourier spectrum, termed the *local wavelet spectrum*, was introduced to quantify this local structure together with an associated estimation theory. A suite of routines for the (unbiased) estimation of the 2D local wavelet spectrum has been implemented as a package in R (see Eckley and Nason (2009) for further details). This software, together with help pages, is available for download via CRAN (<http://cran.r-project.org/>).

We then considered the application of the LS2W modelling approach to texture analysis problems, its potential being contrasted against other recently proposed wavelet-based methods both on a conceptual and applied basis. For many texture classes, such as the pilled images, the LS2W approach was found to achieve classification rates which were comparable with those of the (uncorrected) NDWT – a consequence of the stationary nature of these textures. Clearly the LS2W-approach could also be used for texture synthesis. This is left as an avenue for future research.

The true potential of our model becomes clear when we consider its application to non-stationary texture classification. In this case, the results obtained with an (uncorrected) NDWT approach were found to be inferior to those of the LS2W model. This disparity is due to the latter’s ability to correct for the power leakage which is induced by the redundancy of the NDWT. In future work we hope to demonstrate the potential of the LS2W model to real examples of such structures. One potential application could lie in the area of functional neuroimaging such as fMRI or dynamic positron emission tomography (see for example Hayasaka *et al.* (2004) and Worsley *et al.* (1996)). A particular challenge in this area is the successful identification of localised changes in cerebral activation. Additionally, the LS2W approach could be applied to spatial boundary detection, also referred to as “wombling”. Briefly, the primary inference challenges within wombling are model estimation, spatial prediction and assessment of the estimated spatial surface to detect either (i) physical landmark features or (ii) partition the region into disjoint sets. The latter application is akin to texture segmentation. See Banerjee and Gelfand (2006) for a comprehensive overview of this field. Due to its localised structure, we believe that the LS2W model could contribute to this field either as an intermediate step to the feature detection task or by providing a statistically rigorous framework which can be used for the partitioning of regions (see Csillag and Kabos (2002) for an example of existing work in this area using the DWT).

There are connections with our model (with Gaussian innovations) and Gaussian Markov Random Field (GMRF) models. Indeed, our model can be seen as a hybridization of a GMRF with a multiscale structure, see Rue and Held (2005).

The locally stationary two dimensional process model which we have proposed focuses on analysing the covariance structure on regular grids of size $2^m \times 2^n$. Clearly, it is desirable from a practical perspective to extend such an approach to more general structures, including those with missing observations and/or unevenly spaced locations. The NDWT does not readily lend itself to such extensions, thus alternative approaches, such as lifting transforms, may need to be considered.

Of course, other 2D wavelet decompositions are possible and sometimes preferable. Thus the possibility of deriving alternative model forms arises. In particular, in the future

we are interested in considering anisotropic wavelet bases which might be more suitable for process with differing amounts of smoothness in different directions. Moreover, whilst this model is invariant under translations, it is not rotationally invariant (RI). We therefore hope that *formal* statistical modelling frameworks based on RI-transforms, such as those proposed by Simoncelli and Freeman (1995), will be addressed by future research. It would also be interesting to develop a parallel theory for 2-D LSF and to create associated texture measures. There may be textures more naturally represented by such models. In addition, since NvSK, various relaxations of the Lipschitz conditions which control local stationarity have been proposed. For example, within the time series setting, Fryzlewicz (2003) relaxes the condition on the Lipschitz constants L_j to include time modulated white noise processes, whilst the work of Van Belleghem and Von Sachs (2008) allows for jump discontinuities in the model form.

Finally we turn to texture analysis. Although statistics has devised numerous discrimination and classification schemes which are applied in this field, the issue of obtaining suitable measures from textured images has not yet received much attention in the statistics literature. Our exploratory analyses involving the standard Brodatz collection and (more exacting) industrial collections indicate that none of the approaches considered to date consistently excels. Thus the problem of measure choice is one which is ripe for future research.

7 Acknowledgments

The authors acknowledge the financial support of Unilever Research, are also grateful to the Editor, Associate Editor, several referees and Rebecca Killick for helpful comments.

References

- Abramovich, F., Bailey, T. C., and Sapatinas, T. (2000) Wavelet analysis and its statistical application, *J. R. Statist. Soc. D*, **49**, 1–29.
- Banerjee, S. and Gelfand, A. E. (2006) Bayesian wombling: Curvilinear gradient assessment under spatial process models, *J. Am. Statist. Ass.*, **101**, 1487–1501.
- Bovic, A. C., Clark, M., and Geisler, W. S. (1990) Multichannel texture analysis using localized spatial filters, *IEEE Trans. Patt. Anal. and Mach. Intell.*, **12**, 55–73.
- Chang, T. and Kuo, C.-C. J. (1993) Texture analysis and classification with tree-structured wavelet transform, *IEEE Trans. Im. Proc.*, **2**, 429–441.
- Cressie, N. A. C. (1991) *Statistics for Spatial Data*, Wiley, New York.
- Crouse, M. S., Nowak, R. D., and Baraniuk, R. G. (1998) Wavelet-based statistical signal processing using hidden markov models, *IEEE Trans. Sig. Proc.*, **46**, 886–902.
- Csillag, F. and Kabos, S. (2002) Wavelets, boundaries and the spatial analysis of landscape pattern, *Ecoscience*, **9**, 177–190.
- Dahlhaus, R. (1997) Fitting time series models to nonstationary processes, *Ann. Statist.*, **25**, 1–37.

- Damian, D., Sampson, P. D., and Gutterop, P. (2003) Variance modeling for nonstationary spatial processes with temporal replications, *J. Geophys. Res. Atmos.*, **108**, Art. No. 8778.
- Daubechies, I. (1988) Orthonormal bases of compactly supported wavelets, *Comm. Pure Appl. Math.*, **41**, 909–966.
- Daubechies, I. (1992) *Ten Lectures on Wavelets*, SIAM, Philadelphia.
- Daugman, J. G. (1990) An information-theoretic view of analog representation in striate cortex, in E. L. Schwartz, ed., *Comp. Neurosci.*, pp. 403–423, MIT Press, Cambridge, Mass.
- Davies, L. S. and Mitchie, A. (1980) Edge detection in textures, *Comp. Graph. Im. Proc.*, **12**, 25–39.
- Dunn, D. and Higgins, W. E. (1995) Optimal gabor filters for texture segmentation, *IEEE Trans. Im. Proc.*, **4**, 947–964.
- Eckley, I. and Nason, G. (2009) Ls2w: Locally stationary wavelet fields in R, In preparation.
- Eckley, I. A. (2001) *Wavelet methods for time series and spatial data*, Ph.D. thesis, University of Bristol.
- Eckley, I. A. and Nason, G. P. (2005) Efficient computation of the inner-product matrix of discrete autocorrelation wavelets, *Statistics and Computing*, **15**, 83–92.
- Eckley, I. A., Nason, G. P., and Treloar, R. L. (2008) Technical appendix to lsw fields with application to the modelling and analysis of image texture, Tech. rep., Statistics Group, Lancaster University.
- Field, D. J. (1999) Wavelets, vision and the statistics of natural scenes, *Phil. Trans. Roy. Soc. London (A)*, **357**, 2527–2542.
- Fryzlewicz, P. (2003) *Wavelet Techniques for Time Series and Poisson Data*, Ph.D. thesis, Department of Mathematics, University of Bristol.
- Haas, T. (1990) Lognormal and moving window methods of estimating acid decomposition, *J. Am. Statist. Ass.*, **85**, 950–963.
- Haining, R. P. (1978) The moving average model for spatial interaction, *Trans. Inst. Brit. Geog.*, **3**, 202–225.
- Haralick, R. M. (1979) Statistical and structural approaches to texture, *Proc. IEEE*, **67**, 786–804.
- Haralick, R. M. (1986) Statistical image texture analysis, in T. Young and K. S. Fu, eds., *The Handbook of Pattern Recognition and Image Processing*, pp. 247–279, Academic Press, Orlando, FL.
- Haralick, R. M., Shanmugam, K., and Dinstein, I. (1973) Textural features for image classification, *IEEE Trans. Syst., Man Cyber.*, **SMC-3**, 610–621.
- Hayasaka, S., Phan, K. L., Liberzon, I., Worsley, K. J., and Nichols, T. E. (2004) Nonstationary cluster-size inference with random field and permutation methods, *NeuroImage*, **22**, 676–687.
- Higdon, D., Swall, J., and Kern, J. (1999) Non-stationary spatial modeling, in J. M. B. *et al.* ed., *Bayesian Statistics 6*, pp. 761–768, Oxford University Press, Oxford.
- Kingsbury, N. G. (1999) Image processing with complex wavelets, *Phil. Trans. Roy. Soc. London (A)*, **357**, 2543–2560.
- Laine, A. and Fan, J. (1993) Texture classification by wavelet packet signatures, *IEEE*

- Trans. Patt. Anal. and Mach. Intell.*, **15**, 1186–1191.
- Le, N. D. and Zidek, J. V. (1992) Interpolation with uncertain spatial covariance, *J. Multi. Anal.*, **43**, 351–374.
- Le, N. D., Sun, W., and Zidek, J. V. (1997) Bayesian multivariate spatial interpolation with data missing by design, *J. R. Statist. Soc. B*, **59**, 501–510.
- Loader, C. and Switzer, P. (1992) Spatial covariance estimation for monitoring data, in A. T. Walden and P. Guttorp, eds., *Statistics in the Environmental and Earth Sciences*, pp. 52–70, Arnold, London.
- Mallat, S. G. (1989) A theory for multiresolution signal decomposition: the wavelet representation., *IEEE Trans. Patt. Anal. and Mach. Intell.*, **11**, 674–693.
- Mondal, D. and Percival, D. B. (2008) M-estimation of wavelet variance, In submission.
- Moore, M. (1988) Spatial linear processes, *Comm. Stat.: Stoch. Mod.*, **4**, 45–75.
- Nason, G. P. (2008) *Wavelet Methods in Statistics with R*, Springer.
- Nason, G. P. and Silverman, B. W. (1995) The stationary wavelet transform and some statistical applications, in A. Antoniadis and G. Oppenheim, eds., *Wavelets and Statistics*, number 103 in Lecture Notes in Statistics, pp. 281–300, Springer-Verlag, New York.
- Nason, G. P., von Sachs, R., and Kroisandt, G. (2000) Wavelet processes and adaptive estimation of the evolutionary wavelet spectrum, *J. R. Statist. Soc. B*, **62**, 271–292.
- Nychka, D., Wikle, C., and Royle, A. J. (2002) Multiresolution models for nonstationary spatial covariance functions, *Stat. Mod.*, **2**, 315–331.
- Petrou, M. (2006) *Image Processing: Dealing with Texture*, Wiley.
- Portilla, J. and Simoncelli, E. P. (2000) A parametric texture model based on joint statistics of complex wavelet coefficients, *Int. J. Comp. Vis.*, **40**, 49–70.
- Priestley, M. B. (1981) *Spectral Analysis and Time Series*, Academic Press, London.
- Randen, T. and Husøy, J. H. (1999) Filtering for texture classification: a comparative study, *IEEE Trans. Patt. Anal. and Mach. Intell.*, **21**, 291–310.
- Reed, T. R. and du Buf, J. M. H. (1992) A review of recent texture segmentation and feature extraction techniques, *CVGIP: Image Understanding*, **57**, 359–372.
- Reed, T. R. and Wechsler, H. (1990) Segmentation of textured images and gestalt organization using spatial/spatial-frequency representations, *IEEE Trans. Patt. Anal. and Mach. Intell.*, **12**, 1–12.
- Romberg, J. K., Choi, H., and Baraniuk, R. G. (2001) Bayesian tree-structured image modeling using wavelet-domain hidden Markov models, *IEEE Trans. Im. Proc.*, **10**, 1056–1068.
- Rue, H. and Held, L. (2005) *Gaussian Markov Random Fields: Theory and Applications*, Chapman and Hall.
- Saito, N. and Coifman, R. R. (1995) Local discriminant bases and their applications, *J. Math. Im. Vis.*, **5**, 337–358.
- Sampson, P. D. and Guttorp, P. (1992) Nonparametric estimation of nonstationary spatial covariance structure, *J. Am. Statist. Ass.*, **87**, 108–119.
- Schmidt, A. M. and O’Hagan, A. (2003) Bayesian inference for non-stationary spatial scovariance structure via spatial deformations, *J. R. Statist. Soc. B*, **65**, 743–758.
- Simoncelli, E. P. and Freeman, W. T. (1995) The steerable pyramid: a flexible architecture

- for multi-scale derivative computation, in *Proceedings of the IEEE International Conference on Image Processing, Washington, DC*, volume 3, pp. 444–447, IEEE Computer Society, Los Alamitos, CA.
- Tomita, F. and Tsuji, S. (1990) *Computer Analysis of Visual Textures*, Kluwer, Dordrecht.
- Tuceryan, M. and Jain, A. K. (1999) Texture analysis, in C. H. Chen, L. F. Pau, and P. Wang, eds., *Handbook of Pattern Recognition and Computer Vision (2nd ed)*, pp. 207–248, World Scientific, Singapore.
- Turner, M. R. (1986) Texture discrimination by Gabor functions, *Bio. Cyb.*, **55**, 71–82.
- Unser, M. (1986) Local linear transforms for texture measurements, *Sig. Proc.*, **11**, 61–79.
- Unser, M. (1995) Texture classification and segmentation using wavelet frames, *IEEE Trans. Im. Proc.*, **4**, 1549–1560.
- Van Belleghem, S. and von Sachs, R. (2008) Locally adaptive estimation of evolutionary wavelet spectra, *Ann. Statist.*, **36**, 1879–1024.
- Van de Wouwer, G., Scheunders, P., and Van Dyck, P. (1999) Statistical texture characterization from discrete wavelet representation, *IEEE Trans. Im. Proc.*, **8**, 592–598.
- Venkatachalam, V., Choi, H., and Baraniuk, R. G. (2000) Multiscale SAR image segmentation using wavelet-domain hidden Markov tree models, in *Proc. SPIE 14th Int. Symp. on Aerospace/Defense Sensing, Sim. and Controls, Alg. for Synthetic Aperture Radar Imagery VII, Orlando, FL*, volume 4053, pp. 110–120, SPIE, Bellingham, WA.
- Vidakovic, B. (1999) *Statistical Modelling by Wavelets*, Wiley, New York.
- von Sachs, R., Nason, G. P., and Kroisandt, G. (1997) Adaptive estimation of the evolutionary wavelet spectrum, Tech. Rep. 516, Dept. Statistics, Stanford University.
- Weszka, J. S., Dyer, C., and Rosenfeld, A. (1976) A comparative study of texture measures for terrain classification, *IEEE Trans. Syst., Man Cyber.*, **6**, 269–285.
- Worsley, K. J., Marrett, S., Neelin, P., Vandal, A. C., Friston, K. J., and Evans, A. C. (1996) A unified statistical approach for determining significant signals in images of cerebral activation, *Human Brain Mapping*, **4**, 58–73.